

Toward Evolvable Hardware Chips: Experiments with a Programmable Transistor Array

Adrian Stoica
Jet Propulsion Laboratory
California Institute of Technology
adrian.stoica@jpl.nasa.gov

Abstract

Evolvable Hardware is reconfigurable hardware that self-configures under the control of an evolutionary algorithm. The search for a hardware configuration can be performed using software models or, faster and more accurate, directly in reconfigurable hardware. Several experiments have demonstrated the possibility to automatically synthesize both digital and analog circuits. The paper introduces an approach to automated synthesis of CMOS circuits, based on evolution on a Programmable Transistor Array (PTA). The approach is illustrated with a software experiment showing evolutionary synthesis of a circuit with a desired DC characteristic. A hardware implementation of a test PTA chip is then described, and the same evolutionary experiment is performed on the chip demonstrating circuit synthesis/self-configuration directly in hardware.

*adaptive hardware
programmable devices*

1. Introduction

Evolvable hardware is reconfigurable hardware that self-configures under the control of an evolutionary algorithm. The search for a hardware configuration can be made in software and the final solution can be downloaded to hardware. Alternatively, evolution in hardware (directly on the chip), can speed-up the search for a solution circuit by a few orders of magnitude compared to evolution in software simulations. Moreover, since the software simulation relies on models of physical hardware with limited accuracy, a solution evolved in software may behave differently when downloaded in programmable hardware; such mismatches are avoided when evolution takes place directly in hardware.

Hardware evolution is performed through a succession of changes of elementary cell functions and cell inter-connectivity pattern, thus obtaining increasingly fitter configurations until target functionality is reached. As it is the case in nature,

evolution results in individuals that are increasingly more adapted to their environments, and can change themselves to match changes in environments and modifications of their own goals. Evolution in silicon can however be extremely rapid, with millions of generations of "living" circuits evaluated in only a few seconds.

A variety of circuits have been synthesized through evolutionary means. Koza used Genetic Programming (GP) to grow an "embryonic" circuit to a circuit that satisfies desired requirements [1]. This approach was used for evolving a variety of circuits, including filters and computational circuits. Koza's evolutions were performed in simulations, without concern of a physical implementation, but rather as a proof-of-concept that evolution can lead to designs that compete or even exceed in performance the human designs. No analog programmable devices exist that would support the implementation of the resulting design (but, in principle, one can test their validity in circuits built from discrete components, or in an ASIC), and thus intrinsic evolution was not possible. An alternative encoding technique for analog circuit synthesis, which has the advantage of reduced computational load was used in [2] for automated filter design.

On-chip evolution was demonstrated by Thompson [3]. Thompson used an FPGA as the programmable (digital) device, and a Genetic Algorithm (GA) as the evolutionary mechanism to configure a frequency discriminator from the digital gates available on a small part of the FPGA. Although evolution used the circuitry prepared to implement logic gates, the functionality was obtained exploiting more the underlying physical phenomena at transistor level.

In particular, it is interesting to evolve circuits based on CMOS transistors. CMOS Transistors are the elementary building block of the majority of current microelectronics and addressing evolution at this low level allows most flexibility for synthesizing analog, digital and mixed signal designs. Although for many functions it is easier to synthesize based on higher-level

dedicated blocks, the lessons learned in synthesizing at this level can be extended to evolution of circuits systems made of other devices and materials/structures. An important part of our activity is developing dedicated hardware capable of evolution of both analog and digital circuits, directly on the chip.

This paper proposes a Programmable Transistor Array (PTA) as a platform for experiments in evolutionary synthesis of electronic circuits. On-chip evolutionary experiments with the PTA are expected to lead to design guides for a true stand-alone evolvable chip. An evolution on a simulated PTA illustrates the feasibility of automated synthesis. A chip was designed and fabricated to validate the results in real hardware.

The paper is organized as follows: Section 2 describes the principles of evolutionary synthesis of electronic circuits. Section 3 discusses a GA that acts as the evolutionary self-configuration mechanism and describes an evolutionary design tool developed around a parallel GA package and a circuit simulator. Section 4 proposes a Programmable Transistor Array as an experimental platform for evolutionary synthesis of both analog and digital CMOS circuits. Section 5 describes an experiment in which a CMOS circuit with a Gaussian I-V imposed characteristic was synthesized by evolution. Section 6 discusses hardware aspects related to the implementation of the PTA on a 0.5 micron CMOS test chip. Section 7 presents the evolution directly on the PTA chip and compares the software and hardware experiments. Section 8 presents some lessons learned from the experiments, while Section 9 presents the conclusions of the paper.

2. Principles of evolutionary synthesis of electronic circuits

This section describes the principles of evolutionary synthesis of electronic circuits, and highlights some important results in the field. The idea behind evolutionary synthesis, or EHW is to employ a search/optimization algorithm that operates in the space of all possible circuits and determines solution circuits with desired functional response [4], [5]. Most experiments were performed using evolutionary algorithms such as GA and GP. The genetic search is tightly coupled with a coded representation for the circuits. Each circuit gets associated a "genetic code" or *chromosome*; the simplest representation of a chromosome is a binary string, a succession of 0s and 1s that encode a circuit. Synthesis is the search in the chromosome space for the solution corresponding to a circuit with a desired functional response. The genetic search follows a "generate and test" strategy: a population of candidate solutions is maintained at each time; the corresponding circuits are evaluated and the best candidates are selected and reproduced in a subsequent generation, until a performance goal is reached. Circuit evaluation can be done on software models using circuit simulators, in which case evolution is called *extrinsic* evolution, or directly in reconfigurable hardware, in which case it is called *intrinsic* evolution.

The main steps of evolutionary synthesis are illustrated in Figure 1. First, a population of chromosomes is randomly generated. The chromosomes are converted into circuit models (for extrinsic EHW) or control bitstrings downloaded to programmable hardware




Figure 1: Evolutionary synthesis of electronic hardware

3. Details of the Evolutionary Algorithm and its software implementation within a design tool

A variety of Evolutionary Algorithms (including GA and GP) have been used successfully for evolution of circuits. A GA was chosen here because (a) previous work has demonstrated its efficiency in evolutionary circuit synthesis, (b) the mechanism is simple to understand and implement, (c) public domain software exists and saves development time, and (d) the focus was on the reconfigurable hardware and not on the reconfiguration mechanism. It is likely that more intelligence can be inserted into the search mechanism. A simple block diagram of operations taking place in a GA is illustrated in Figure 2.



testing architectures of reconfigurable HW and demonstrating evolution on them before the fabrication of a dedicated reconfigurable chip. The tool can also be used in hardware-software co-design. In its current implementation the tool uses the public domain Parallel Genetic Algorithm package PGAPack and a public domain version of SPICE 3F5 as circuit simulator. An interface code links the GA with the simulator where potential designs are evaluated, while a GUI allows easy problem formulation and visualization of results. Each generation the GA produces a new population of binary chromosomes, which get converted into voltages in netlists that describe candidate circuit designs, netlists further simulated by SPICE.



This section introduces evolution of CMOS circuits based on Programmable Transistor Arrays, describing a design for hardware reconfigurable at transistor level. The PTA allows synthesis of analog, digital and mixed-signal circuits, being a more suitable platform for synthesis of analog circuitry than existing FPGAs or FPAAs, extending the work on evolving simulated circuits to evolving analog circuits directly on the chip.

The PTA idea was introduced in [6], and expanded in [7]. The proposed PTA is an array of transistors interconnected by programmable switches. The status of the switches (On or Off) determines a circuit topology and consequently a specific response. Thus, the topology can be considered as a function of switch states, and can be represented by a binary sequence, such as “1011...”, where by convention one can assign 1 to a switch turned On and 0 to a switch turned Off. The PTA is a modular

architecture, in which modules can be cascaded to determine a more complicated circuit topology. Figure 4 illustrates an example of a PTA module consisting of 8 transistors and 24 programmable switches. In this example the transistors P1-P4 are PMOS and N5-N8 are NMOS, and the switch based-connections are in sufficient number to allow a majority of meaningful topologies for the given transistors arrangement, and yet less than the total number of possible connections.




Figure 4: Module of the Programmable Transistor Array

Programming the switches On and Off determines a circuit for which the effects of non-zero, finite impedance of the switches can be neglected in the first approximation. An example of a circuit drawn with this simplification is given in Figure 5.




Figure 5: Schematic of a simple circuit implemented on the PTA module (with finite resistance of Off switches as dotted lines on the right figure)

The left drawing illustrates the ideal circuit, the right drawing shows with dotted lines the finite resistance of open switches. A power supply, input signals and a biasing current source have been added.

In this implementation four layers of transistors (two PMOS and two NMOS) were chosen, but this can be increased, for example to 6 or 8. On the "horizontal" direction the PTA architecture allows implementing bigger circuits by cascading PTA modules. A simple expansion would be by connecting two adjacent modules with a set of programmable connections. One such expansion with 24 connections between two modules (and thus a total of 72 programmable elements) was simulated. Although further research is needed before conclusive remarks can be made, the cascaded ensemble of two modules has shown rich behavior, and was able to evolve solution circuits to the experiments presented in the next section. It is likely that a cell-based architecture with the same rich capabilities as found in FPGAs (and possibly sharing architectural design ideas) will be needed.

One important question is "how do we know the size of the circuit that the evolution would synthesize?" In the simulated experiments performed so far a choice was to use one or two modules, but for arbitrary functions, without a prior knowledge of a human-designed circuit, we may not have a clear estimate of what is a good size (is it one, two or ten modules?). If we choose fewer modules than necessary we perform a search in a space where there is no solution. On the other hand, a too big search space complicates the evolutionary search.

One possible solution is to use many internal testpoints on the PTA as possible outputs and narrow down the selection based on a distance of the response probed there and target response. This can be individuals in the same population with different sizes (chromosome length) or simply routed outputs and parallel evaluation of many circuits. A part of the individual genetic code would indicate where is the output probed.

5. An experiment in evolutionary CMOS circuit synthesis on a simulated PTA

This section details the evolution of a circuit with a Gaussian I-V DC response. The evolutionary synthesis approach illustrated in Figure 1 was applied to the model of PTA illustrated in Figure 4.

Evolutionary synthesis of a computational circuit was chosen to illustrate the approach. The goal of evolution was to synthesize a circuit which exhibits a Gaussian I-V characteristic. In a previous experiment [8] the circuit topology was fixed and the search search/optimization

addressed transistor parameters (channel length and width); such evolution proved quite simple. The search for a topology turned out to be a much harder problem and several architectures were unsuccessfully attempted before the PTA was conceived. In the PTA case, the transistor parameters were kept fixed and the search was performed for the 24 binary parameters characterizing switches status. An important role was the correct specification of the fitness function, for which a weighted combination of parameters x_1, \dots, x_7 in Figure 6 was used.

The evolution was simulated on a Caltech supercomputer (HP-Exemplar), using the Evolutionary Design Tool. Successful evolution was demonstrated on multiple runs with populations between 50 and 512, evolving for 50 or 100 generations. The execution time depends on the above variables and on the number of processors used (usually 64 out of the 256 available), averaging around 20 minutes (the same evolutions took about 2 days on a SUN SPARC 10). In some runs the solution circuit shown in Figure 5 (human designed) was rediscovered by evolution.




Figure 6: Parameters used for the specification of the fitness function. $Fitness = f(x_1, \dots, x_7)$

Other solutions found include the circuits illustrated in Figure 8, which produce the first two responses in Figure 7; some other responses from the same generation are illustrated in Figure 7 for comparison. It is interesting to analyze in more detail the unusual solutions found by evolution. Circuits like those illustrated in Figure 8 resulted from evolutionary synthesis are very similar (under certain test conditions) to that of the circuit shown in Figure 5. Thicker dotted lines show connections that existed in the circuit in Figure 5, but are missing in the circuits in Figure 8. As it is easy to observe these circuits are outside normal design practices, e.g., the transistors P2, P4 and N8 on the left circuit in Figure 8 and P2 have floating gates. The reality is that the switches have a big, but finite, resistance in the Off state (\sim MOhms or

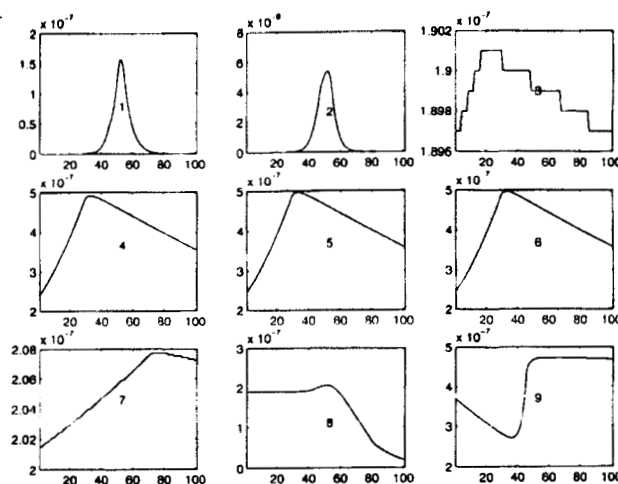


Figure 7: Best circuit responses in a simulated evolution

GOhms) and a non-zero resistance/impedance in the On state (\sim tens of Ohms). *One observation from here is that while the effects of non-perfect switches may be negligible in a first approximation for many digital circuits, such effects may fundamentally affect analog programmable circuits.*




Figure 8: Circuits obtained by evolution; their design is unusual for common practice

6. Hardware Implementation

Successful evolution on simulated PTA encouraged the development of a test chip implementing the PTA architecture. The chip would offer an estimate on how reliable is evolution on SW models. More importantly, evolution of the circuit directly on the chip becomes possible, and at an expected accelerated pace of over 100 times compared to the simulation (estimated \sim 5 seconds compared to \sim 20 minutes on the supercomputer for the experiment described). As in the experiments performed in simulations the size of transistors was fixed. The

programmable switches were implemented with transistors, acting as simple T-gate switches. There were several considerations for this choice:

- The switches has to pass analog signals
- The resistance of the switches needed to be variable between low (~tens/hundreds of ohms) and high (in excess of tens and hundreds of MOhms and above).
- Intermediate resistance was necessary (for experiments that will be described elsewhere) but linearity ($R=R(V_{gate-control})$) was not important

One should mention that the analog gradual switches act in circuit evolution very much like resistive weights in a neural network implementation.

Each chip implements one PTA module. To offer sufficient flexibility the chip has all transistor terminals (except those connected to power and ground) connected via switches to expansion terminals. A crossbar switch array could be used, but in the initial experiments some connections of choice will be wired. The switches will be controlled in tandem to ensure a connection between the terminals of the two modules.

The chip was fabricated as a Tiny Chip through MOSIS, using 0.5-micron CMOS technology. The test board with four chips mounted on it is illustrated in Figure 9.




Figure 9: A test board with four PTA Chips

7. Evolution on the PTA chip

The same evolutionary experiment, aiming at the synthesis of a DC circuit with a Gaussian response, was performed in hardware on the PTA chips, (the GA was implemented in LabView). Four chips were programmed in parallel with bit-string configurations corresponding to four individuals of a population of 100; then, the next four were programmed, and so on until all 100 in one generation were tested. As in simulation, evolution led to "Gaussian" circuit solutions within 200-300 generations. The response of four mutants is illustrated in the screen capture shown in Figure 10 (LabView display of the signals captured by the data acquisition boards). Notice the "mutations" in the genetic code of the solutions (vertical chromosomes R24 - R1 reading from top to bottom - these correspond to switches S24 - S1 in Figure 4) compared with the generic solution.




Figure 10: The "Gaussian" response of four "mutants" and their "genetic code" compared to the generic solution

8. Lessons learned

- An interesting observation was that, *other than the "correct" human-designed solution rediscovered by evolution*, the solutions evolved on the PTA chip are different than those evolved in simulations. (At least the few of them that were tested; additional circuits solutions may exist that lead to the same response both in the Spice simulation and programmed on the chip). It would thus appear that different effects are exploited to lead to solutions in the model and in the silicon implementation. More precisely, the circuit solutions evolved in simulations (with Spice resistive models for On/Off switches) did not prove to be solutions when programmed on the PTA chip, and vice-versa, the configuration solutions evolved directly on the PTA chip (e.g. those in Figure 10) did not simulate as Gaussians. (Further experiments using more accurate models of the PTA silicon implementation are in progress). Thus, it appears justifiable to express reserve on the validity of a solution obtained by "extrinsic" evolution of analog circuits until is verified in hardware (at least for particular PTA discussed here and with the limited accuracy model used).
- The original intent was to speed-up the evolution from ~20 minutes on the supercomputer to about 5 seconds on the PTA chip (reducing the evaluation of a circuit to ~0.25ms). At this moment, LabView (running on a 300 MHz Pentium) presents some communication bottlenecks that only allowed reaching about the same evolution time as on supercomputer. In the quest for faster circuit

evaluation on the chip a further limitation was however noticed, ignored when running Spice DC analysis only: the circuits have own frequency response and there are limits of possible speed-up for which the response is the same as in DC/low frequency. The output of the Gaussian circuit on the PTA starts attenuation when the input ramp signals exceed 1kHz, meaning that that no more than 1000 circuits per second could be reliably evaluated. Even though this may be an artifact of the particular PTA design and load choice, it appears natural that evaluating the circuits at a different frequency than that of intended functioning may introduce errors. Evaluation in parallel is an alternative speed-up technique, and at least in the experiments with the PTA chips no significant differences were noted between the instantiations of the same circuit on different chips.

9. Conclusion

Automatic synthesis/self-configuration of analog circuits was demonstrated on an experimental CMOS chip implementing a Programmable Transistor Array architecture proposed as reconfigurable hardware platform for evolutionary synthesis experiments. The experiments bring further testimony to the feasibility of using evolutionary algorithms for automated synthesis of electronic circuits. A comparison of the simulated and on-chip experiments indicates limitations of the extrinsic evolutionary method; the solutions obtained in simulations were not validated when programmed on the chip. However, different solutions have evolved on the chip, and proved robust when transferred to other chips from the same fabrication lot.

10. Acknowledgements

The research described in this paper was performed at the Center for Integrated Space Microsystems, Jet Propulsion Laboratory, California Institute of Technology and was sponsored by the National Aeronautics and Space Administration.

Several people contributed to the development presented here. Two JPL summer intern students deserve special acknowledgement and thanks for their dedicated effort: Carlos Salazar-Lazaro, from Rensselaer Polytechnic Institute, who wrote most of the simulation software, and Wei-te Li, from University of Washington, who did most of the chip design work. Raoul Tawel contributed to chip design, and Ken Hayworth and Didier Keymeulen

contributed to testing the chip and evolution directly on the chip. The author also wishes to thank the reviewers of this paper for their useful suggestions.

9. References

- [1] J. Koza, F.H. Bennett, D. Andre, and M.A. Keane, "Automated WYWIWYG design of both the topology and component values of analog electrical circuits using genetic programming", *Proceedings of Genetic Programming Conference*, Stanford, 1996, pp. 28-31.
- [2] J. Lohn, J. and S. Colombano, "Automated Analog Circuit Synthesis using a linear representation", M. Sipper, D. Mange and A. Perez-Urbe (Eds.) *Evolvable Systems: From Biology to Hardware*, Springer-Verlag Lecture Notes in Computer Science Berlin 1998, pp. 125-133
- [3] A. Thompson, An evolved circuit, intrinsic in silicon, entwined in physics. In *International Conference on Evolvable Systems*. Springer-Verlag Lecture Notes in Computer Science, 1996, pp. 390-405.
- [4] De Garis, H. Evolvable Hardware: Genetic Programming of a Darwin Machine. *Int. Conf. on Artificial Neural Networks and Genetic Algorithms*, Innsbruck, Austria, Springer Verlag, 1993
- [5] Higuchi T. et al, Evolvable Hardware with Genetic Learning, in *Proc. of Simulated Adaptive Behavior*, MIT Press, 1993
- [6] Stoica, A., "Reconfigurable Transistor Array for Evolvable Hardware", Caltech/JPL Novel Technology Report, July 26, 1996
- [7] Stoica, A. and Salazar-Lazaro, C. "Evolutionary technique for automated synthesis of electronic circuits" Caltech/JPL Novel Technology Report, Sept. 4, 1998
- [8] Stoica, A., On hardware evolvability and levels of granularity. In *International Conference on Intelligent Systems and Semiotics*, NIST Gaithersburg VA, Sept, 1997